Shape Distributions of Nonlinear Dynamical Systems for Video-based Inference

Vinay Venkataraman, *Student Member, IEEE*, and Pavan Turaga, *Senior Member, IEEE*

Abstract—This paper presents a shape-theoretic framework for dynamical analysis of nonlinear dynamical systems which appear frequently in several video-based inference tasks. Traditional approaches to dynamical modeling have included linear and nonlinear methods with their respective drawbacks. A novel approach we propose is the use of descriptors of the shape of the dynamical attractor as a feature representation of nature of dynamics. The proposed framework has two main advantages over traditional approaches: a) representation of the dynamical system is derived directly from the observational data, without any inherent assumptions, and b) the proposed features show stability under different time-series lengths where traditional dynamical invariants fail. We illustrate our idea using nonlinear dynamical models such as Lorenz and Rossler systems, where our feature representations (shape distribution) support our hypothesis that the local shape of the reconstructed phase space can be used as a discriminative feature. Our experimental analyses on these models also indicate that the proposed framework show stability for different time-series lengths, which is useful when the available number of samples are small/variable. The specific applications of interest in this paper are: 1) activity recognition using motion capture and RGBD sensors, 2) activity quality assessment for applications in stroke rehabilitation, and 3) dynamical scene classification. We provide experimental validation through action and gesture recognition experiments on motion capture and Kinect datasets. In all these scenarios, we show experimental evidence of the favorable properties of the proposed representation.

Index Terms—Action modeling, largest Lyapunov exponent, chaos theory, shape distribution, action and gesture recognition, movement quality assessment, dynamical scene analysis.

arXiv:1601.07471v1 [cs.CV] 27 Jan 2016

1 Introduction

Dynamical modeling methods for understanding signals from various sensing platforms have been the cornerstone of many applications in the computer vision community, such as human activity analysis [2] and dynamical natural scene recognition [3]. Recent advances in sensing platforms like motion capture systems and the Kinect have opened doors to several applications including home-based health monitoring, gaming and entertainment. Take for instance, the task of developing algorithms for understanding the dynamics in human activities. This problem is non-trivial due to the complexity of natural human movement, which is a result of interactions between multiple body joints having high degrees of freedom. In addition, the task of recognizing human actions is challenging due to several factors including inter-class similarities between actions (e.g., running and walking), intra-class variations due to multiple strategies for an action (e.g., dance) and inter-subject variations. Natural human movements (such as walking, running) are composed of periodic action sequences in the form of repetitions, with some variability [4]. These inherent attributes of human movement (periodicity with variability) descriptive of a complex nonlinear chaotic system has motivated researchers to employ tools from nonlinear dynamical systems theory to model human movement [4], [5], [6], [7], [8], [9], [10], [11]. Dynamical modeling of spatio-temporal evolution of human activities are traditionally accomplished by defining a state space and learning a function that maps the current state to the next state [12], [13]. A recent alternate approach has attempted to derive a representation for the dynamical system directly from the observation data using tools from chaos theory [5]. The main idea here is that, by using a top-down approach of dynamical modeling, one would only approximate the true-dynamics of the system with attempts to fit a model to the observational data. Whereas, in the bottom-up approach [5], the dynamical system parameters such as the number of independent variables, degrees of freedom and other unknown parameters are estimated from the data. Such an approach can be seen as a generalized representation without any strong assumptions, suitable for analyzing a wide range of dynamical phenomenon.

• *V. Venkataraman and P. Turaga are with the School of Electrical, Computer and Energy Engineering and School of Arts, Media and Engineering, Arizona State University, Tempe, USA.*

E-mail: vvenka18@asu.edu and pturaga@asu.edu

2 Related Work

Several approaches have been proposed in literature for modeling the dynamics in an observed time-series and we list the prior works in the specific applications of interest in a) activity recognition, b) activity quality assessment, and c) natural scene recognition.

2.1 Classical Dynamical Invariants

The largest Lyapunov exponent is a widely used dynamical invariant (measure of chaos), which quantifies the rate of divergence of initially closely-spaced trajectories [3], [5]. A practical method for estimating the largest Lyapunov exponent from an observational time-series was first proposed by Wolf *et al*. [14]. Several other approaches were also proposed in literature to *quantify* chaos [15], [16], [17], which were found to suffer from at least one of these drawbacks: (a) unreliable for small datasets, (b) computationally intensive, (c) relatively difficult to implement [18]. An improved method for estimation of the largest Lyapunov exponent to overcome the above mentioned drawbacks was later proposed by Rosenstein *et al*. [18]. However, experimental results on nonlinear dynamical models have shown that the suggested number of data samples for accurate estimation of the largest Lyapunov exponent is 10^m (where m is the embedding dimension) [18], [19]. In recent years, these methods have been applied to model various visual dynamical phenomenon such as video-based recognition of human activities [5] as well as recognition of dynamical scenes [3]. However, when one needs to make inferences from short videos, or for instance when the activity of interest lasts only a few seconds, the classical approaches have significant drawbacks. While quantification of chaos using the largest Lyapunov exponent have been used to monitor varying chaos levels (level of complexity of the system) for recognition or prediction purposes [20], experimental studies for modeling human activities have not reported any evidence for different levels of chaos in human activities. Hence, we believe that a representation for level of chaos may not be a suitable approach to model human activities. In this paper, we propose an alternative approach to model human activities by extracting dynamical features representative of the *shape* of the reconstructed phase space instead of quantifying chaos. We also demonstrate through experiments that the framework for estimation of dynamical features show stability across different time-series lengths and compare the performance with traditional chaotic invariants.

2.2 Activity Recognition

Human activity analysis has attracted the attention of many researchers providing extensive literature on the subject. A detailed review of the approaches in literature for modeling and recognition of human activities are discussed in [2], [21]. Since our present work is related to non-parametric approaches for dynamical system analysis for action modeling, we restrict our discussion to related methods.

Human actions have been modeled using dynamical system theory in computer vision [5], [13] and biomechanics [4], [7], [8]. Differential equations can be used to model such a system, which requires access to all independent variables of the system. This approach would facilitate an understanding of the system behavior and also allow for the prediction of future states using present and past state information. However, this is not realizable in practice, as it is extremely hard to determine the independent variables and the interactions governing the dynamics of human actions.

Dynamical modeling of human actions can be broadly categorized into parametric and nonparametric methods. Furthermore, human actions have been modeled with the assumption that the underlying dynamical system is linear [13] or nonlinear [5], [12]. In parametric modeling approaches, the dynamics of a system is represented by imposing a model and learning the model parameters from training data. Hidden Markov Models (HMMs) [22] and Linear Dynamical Systems (LDSs) [23] are the most popular parametric modeling approaches employed for action recognition [24], [25], [26], [27] and gait analysis [13], [28], [29]. Nonlinear parametric modeling approaches like Switching Linear Dynamical Systems (SLDSs) have been utilized to model complex activities composed of sequences of short segments modeled by LDS [30]. While, nonlinear approaches can provide a more accurate model, it is difficult to precisely learn the model parameters. In addition, one would only approximate the true-dynamics of the system with attempts to fit a model to the experimental data. An alternative nonparametric action modeling approach is based on tools from chaos theory, with no assumptions on the underlying dynamical system. Traditional chaotic measures, like the largest Lyapunov exponent, correlation dimension and correlation integral, have been extensively used to model human actions [4], [5], [7], [8]. However, [18] and [19] have shown that these nonlinear dynamical measures need large amounts of data to produce stable results (10^m, where m is the embedding dimension). Junejo *et al*. [6] used a self-similarity matrix, a graphical representation of distinct recurrent behavior of nonlinear dynamical systems, to learn an action descriptor. In this paper, through illustrative examples and experimental validation, we show that our framework works better than traditional chaotic invariants for action modeling.

2.3 Activity Quality for Stroke Rehabilitation

Recently researchers from various backgrounds have shown interest in the development of computational frameworks for quantification of *quality* of movement, for possible applications in health monitoring and rehabilitation [1], [4], [19], [31]. Stroke being the most common neurological disorder, leaves millions disabled every year who are unable to undergo long-term therapy treatment due to insufficient coverage by insurance. Recent directions in rehabilitation research

has been towards development of portable systems for therapy treatment. Traditional quantitative scales such as the Fugl Meyer Test [32] and the Wolf Motor Function Test (WMFT) [33], have proven to be effective in evaluating movement quality. However, these approaches involve visual monitoring which would greatly benefit from the development of an objective computational framework for movement quality assessment. The aim here is to develop standardized methods to describe the level of impairment across subjects. We show the utility of the proposed action modeling framework for quantifying the quality of reaching tasks using a single marker on the wrist, and obtain comparable results to a heavy marker-based setup (14 markers placed on arm, shoulder and torso [31]).

The focus of existing approaches for movement quality assessment has been towards finding typical patterns in kinematics which differ between healthy and impaired subjects. While these approaches are successful in giving an insight into understanding human movement, they fail to utilize the inherent dynamical nature of the movement. Rehabilitation therapies are composed of repetitive movements (e.g., reach to a target) that are strongly periodic with inherent variability. Traditional methods have assumed that this variability arises from noise in the system. However, it is evident that variability is an integral part of repetitive movements due to the availability of multiple strategies for the movement. Also, it is believed that variability produced in human movement is a result of nonlinear interactions and have deterministic origin [4]. Extensive research has been carried out to model this variability using nonlinear dynamical system theory [4], [7], [8]. In this paper, we utilize the action modeling framework for movement quality assessment using a single wrist marker.

2.4 Natural Scene Classification

Natural scene classification has been an active area of research in computer vision with applications in automated image and video understanding. Much research has been focused around scene classification using single still images [34], [35], thereby neglecting dynamical motion information available in videos. Recently, the problem of dynamical modeling of natural scenes was introduced by Shroff *et al.* [3] who utilized tools from chaos theory along with GIST [36], [37] to model the spatio-temporal evolution in natural scenes in an unconstrained setting.

Dynamic texture representation using LDS proposed by Soatto *et al*. have been used to recognize and synthesize dynamic textures such as sea-waves, smoke, traffic [38], [39]. Such low-dimensional models have been used to capture complex natural phenomena. However, experimental results reported in [3] show that these simple models might not be effective for dynamic scene classification in an unconstrained setting. Shroff *et al*. utilized traditional chaotic invariants to model the dynamics and have shown that dynamical attributes augmented with spatial attributes (GIST [37]) can be effectively used for categorization of dynamic scenes [3]. Another recent approach utilized spatio-temporal oriented energy filters for dynamic natural scene classification [40]. In this paper, we test the generality of the proposed action modeling framework for dynamic scene classification application.

Contributions: In this paper, we present a computational framework for analysis of dynamical systems by combining the theoretical concepts of dynamical system analysis and ideas in shape theory. We extract dynamical shape features from the reconstructed phase space in the form of shape distributions to achieve improved results. We show the utility of the proposed framework in action and gesture recognition, movement quality assessment and dynamical scene recognition and evaluate the performance by comparing it with traditional chaotic invariants. We also propose two new shape functions to encode local dynamical evolution as opposed to global shape functions proposed by Osada *et al*. [41].

3 Preliminaries

In this section, we introduce the background necessary to develop an understanding of nonlinear dynamical system analysis and chaos theory for applications in activity analysis, activity quality assessment and natural scene analysis.

3.1 Dynamical System Analysis

Dynamical systems are governed by a set of functions defining the variations in the behavior of the system over time. A dynamical system is termed linear or nonlinear if the function defining the behavior of the system is linear or nonlinear respectively. Dynamical systems can be represented using state variables defining the state of the system at a given time t. A dynamical system is termed deterministic if there exists a unique future state for a given current state and is termed stochastic if the future state is derived from a probability distribution of possible states. Chaos theory is the field of study of such deterministic dynamical systems that show high sensitivity to initial conditions. A chaotic system is a dynamical system with deterministic behavior showing sensitivity to initial conditions.

The states of a chaotic system are generally considered to be in an n-dimensional manifold also called *phase space*. A chaotic system evolves over time in its phase space according to the system variables governing the dynamics. The path traversed by the system over time is called a *trajectory* and the region

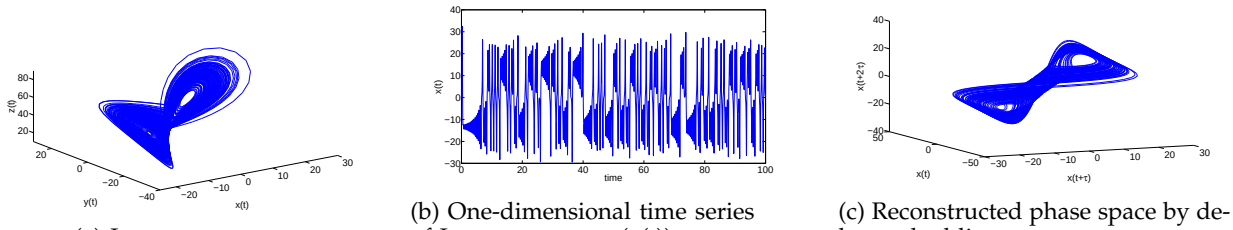

(a) Lorenz attractor

(b) One-dimensional time series of Lorenz attractor (x(t))

(c) Reconstructed phase space by delay embedding

Fig. 1: Phase space reconstruction of Lorenz attractor by delay embedding. (a) shows the 3D view of trajectories of Lorenz attractor with control parameters $\rho = 45.92, \sigma = 16.0$ and $\beta = 4.0$. We can see that trajectories of Lorenz system settle down and are confined within the attractor. The one-dimensional time series (observed) of the Lorenz system is shown in (b). We see that a low-dimensional nonlinear system can generate such complex and chaotic signal. (c) shows the reconstructed phase space from observed time series of the Lorenz system using delay embedding ($\tau = 11$). The above example illustrates that the reconstructed phase space preserves certain topological properties of the original Lorenz attractor.

of the phase space where the trajectories settle down as time approaches infinity is denoted as an *attractor*.

One would intend to have access to all independent variables of the system and their interactions for a complete understanding of the system. In a real world scenario, the data recorded is of low-dimension and is insufficient to model the dynamics of the system. In addition, model-based (parametric) approaches, such as LDS assume an underlying mapping function f to describe the dynamics of the system. It has been established that such approaches may not be suitable for modeling the dynamics of complex systems such as human movements due to the simplifying assumptions [42]. The theory of chaotic systems allows for determining certain invariants of the dynamical system function f without making any assumptions about the system.

3.2 Phase Space Reconstruction

The *phase space* is defined as the space with all possible states of a system [43], [44]. In a deterministic dynamical system that can be mathematically modeled, future states of the system can be determined using present and past state information. However, for applications such as human activity understanding and dynamical scene understanding, the system equations are complex. Furthermore, sensing systems in the real-world do not allow us to observe all variables of the system (e.g., the home-based setting for stroke rehabilitation with single marker on the wrist). To address these problems, we have to employ methods for reconstructing the attractor to obtain a phase space which preserves the important topological properties of the original dynamical system. This process is required to find the mapping function between the one-dimensional observed time series and the m-dimensional attractor, with the assumption that all variables of the system influence one another. The concept of phase space reconstruction was expounded in the embedding theorem proposed by Takens, called Takens' embedding theorem [45] and an example of the procedure is shown in Fig. 1. For a discrete dynamical system with a multidimensional phase space, time-delay vectors (or embedding vectors) are obtained by concatenation of time-delayed samples given by

$$\mathbf{x}_i(n) = [x_i(n), x_i(n+\tau), \cdots, x_i(n+(m-1)\tau)]^T, \quad (1)$$

where 'm' is the embedding dimension and 'τ' is the embedding delay. These parameters should be carefully selected in order to facilitate a good phase space reconstruction. For a sufficiently large 'm', the important topological properties of the unknown multidimensional system are reproduced in the reconstructed phase space [44]. The embedding method has proven to be useful, particularly for time series generated from low-dimensional deterministic dynamical systems, by providing a way to apply theoretical concepts of nonlinear dynamical systems onto observed time series. The embedding theorem does not suggest methods to estimate the optimal values for 'm' and 'τ'. We use the false nearest neighbors [46] approach to estimate m and the first zero crossing of the autocorrelation function [47] to estimate τ. Fig. 1 shows an example of phase space reconstruction from a one-dimensional observed time-series of a Lorenz system.

3.3 Embedding Dimension

The embedding dimension refers to the number of time-delayed samples concatenated to form the time-delay vector (see (1)). The aim here is to estimate an integer embedding dimension which can *unfold* the attractor thereby removing any self-overlaps due to projection of the attractor onto lower dimensional space. Hence, the embedding dimension can be defined as the minimum dimension required to unfold the attractor completely. The false nearest neighbor approach finds this minimum embedding dimension to remove any *false* nearest neighbors (neighbors due to projection onto lower dimension) [44]. Consider a vector in reconstructed phase space in dimension m given by

$$\mathbf{x}(k) = [x(k), x(k+\tau), \cdots, x(k+(m-1)\tau)]^T, \quad (2a)$$

and a nearest neighbor in the phase space given by

$$\mathbf{x}^{NN}(k) = [x^{NN}(k), x^{NN}(k+\tau), \cdots, x^{NN}(k+(m-1)\tau)]^T. \quad (2b)$$

If the vector $\mathbf{x}^{NN}(k)$ is a true neighbor of $\mathbf{x}(k)$, then it should be because of the underlying dynamics. The vector $\mathbf{x}^{NN}(k)$ can be a false neighbor of $\mathbf{x}(k)$ when dimension m is unable to unfold the attractor. Hence, moving to the next dimension $m+1$ may move this false neighbor out of the neighborhood of $\mathbf{x}(k)$. This process of finding false neighbors to every vector $\mathbf{x}_i(k)$ sequentially removes self-overlaps and identifies m where the attractor is completely unfolded. The embedding dimension m suggested by the false nearest neighbor algorithm for exemplar trajectories of human actions was either 3 or 4. We select a constant embedding dimension $m = 3$ to reconstruct all relevant phase space. Even with this fixed value of m, we obtain excellent results as shown in our experiments.

3.4 Embedding Delay

Embedding delay refers to the choice of integer time delay used to construct the time-delay vector. Theoretically, the embedding process allows any value of τ if one has access to infinitely accurate data ([44], chap. 3). Since this is practically impossible, we try to find a value τ which makes the components of the vector $[x(k),\ x(k+\tau),\ x(k+2\tau)]^T$ in the embedding sufficiently independent. A low value of τ makes adjacent components to be correlated and hence they cannot be considered as independent variables. On the other hand, a high value of τ may make the adjacent components uncorrelated (almost independent) and cannot be considered as part of the system that supposedly generated them. The shape of the embedded time series will critically depend on the choice of τ [47]. A good selection of τ should ensure that the data are maximally spread in phase space resulting in smooth phase space reconstruction. We use the first zero-crossing of the autocorrelation function as an estimate of τ as suggested in [47] for strongly periodic data, which is a suitable choice for our experiments.

3.5 Phase Space Reconstruction of the Lorenz Attractor

The Lorenz attractor is the steady state of a nonlinear chaotic system of three coupled nonlinear ordinary differential equations [48] as given below:

$$\dot{x} = \sigma(y - x), \quad (3a)$$

$$\dot{y} = x(\rho - z) - y, \quad (3b)$$

$$\dot{z} = xy - \beta z, \quad (3c)$$

where x, y, z are the state variables and σ, ρ and β are non-negative and dimensionless parameters. These equations were defined by Lorenz in 1963 [43] to represent a simplified model of thermal convection in the lower atmosphere. Lorenz showed that this relatively simple-looking set of equations could have highly erratic dynamics for a range of defined control parameters, for which the dynamics are chaotic. The dynamics of the Lorenz system in the 3-dimensional state space generated from these set of equations is illustrated in Fig. 1(a). Lorenz attractor also illustrates that deterministic nonlinear models of low dimension can produce signal with complex dynamics. Furthermore, Fig. 1 illustrates that it is possible to recreate an approximate attractor generated by a multidimensional system (such as Lorenz) using only a one-dimensional observed time series.

In the next section, we propose dynamical shape feature extraction from reconstructed phase space which is more suitable for action modeling than traditional chaotic invariants. We also show the stability of the proposed dynamical shape features for different time-series lengths using nonlinear dynamical models (Lorenz and Rossler systems).

4 Attractor Shape Distributions

In this section, we present a framework which combines the strong theoretical concepts of nonlinear dynamical analysis and ideas in shape theory to effectively represent the nature of dynamics. From Fig. 2, we see that the '*shape*' of the reconstructed phase space can be seen as a discriminative feature for classification between Run and $Walk$ action classes. Hence, our aim will be to extract feature representations for the shape of the reconstructed phase space. It is important to note here that the process of phase space reconstruction preserves certain topological properties and global shape is not a topological invariant, while local shape properties are. However, our goal here is to suggest a shape-based descriptor (both global and local) which possess sufficient discriminatory properties and robustness.

We consider the attractor as having its own characteristic shape in the high-dimensional phase space. Shape analysis of 3D surfaces is a well-studied problem in the computer vision community. In [41], Osada *et al*. present a method for finding a similarity measure between 3D shapes by computing shape distributions of the 3D surface sampled from the shape function by measuring their global geometric properties. We use the shape distribution of the reconstructed phase space as the dynamical feature representation in our experiments. While the shape distributions was originally proposed to measure similarity between 3D shapes, we believe that shape distributions can be used as feature representations for any n-dimensional phase space. In addition, it is said that any function can be used to extract the shape distribution [41], but we adopt simpler shape functions based on geometric

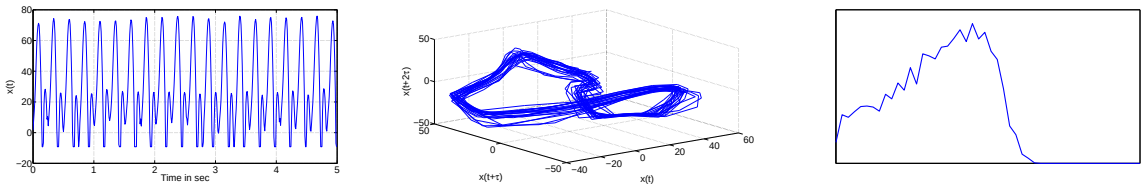

(a) Time series data (b) Reconstructed Phase Space (c) Shape Distribution

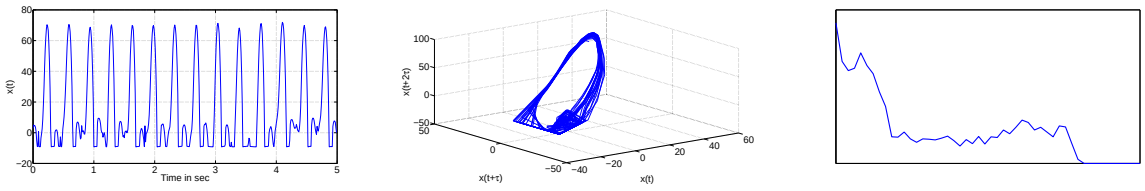

(d) Time series data (e) Reconstructed Phase Space (f) Shape Distribution

Fig. 2: Examples of phase space reconstruction of corresponding time series data of a subject performing *Run* and *Walk* action respectively. The embedding parameters were selected as $m = 3$ and τ as described in section 3.4. This example illustrates that the *shape* of the reconstructed phase space can be seen as a discriminative feature for classification of actions. We use shape distributions proposed by Osada *et al*. [41] as a representation for shape of phase space. (c) and (f) together support our hypothesis that shape distribution (**D2**) can be used for classification of actions.

properties (distance and area) which are listed below:
(a) *Global Shape Functions*:

- **D1**: measures the distance between one fixed point and one random point sampled from the reconstructed phase space. The fixed point is selected as the centroid of the attractor.
- **D2**: measures the distance between two random points in the phase space represented as $||\mathbf{x}_i - \mathbf{x}_j||_2$.
- **D3**: measures the square root of the area of the triangle formed by three random points on the attractor.

For example, the **D2** shape function can be represented as

$$\mathbf{D2}_{ij} = ||\mathbf{x}_i - \mathbf{x}_j||_2, \tag{4}$$

where $\mathbf{x}_i$ and $\mathbf{x}_j$ are points (embedding vectors) in the reconstructed phase space. A set of these distances for randomly chosen embedding vector pairs are computed. From this set, we construct a histogram by counting the number of samples which fall into each of B=50 fixed sized bins to obtain the attractor's shape distribution.

These shape functions encode global geometric properties of the phase space, lacking information about local shape and dynamical evolution in the phase space. While previous investigation shows that global geometric shape function (**D2**) performs sufficiently better than the traditional nonlinear dynamical measures (largest Lyapunov exponent, correlation dimension and correlation integral) [1], we hypothesize that a shape function which encodes local geometry and dynamical evolution information of phase space should improve the performance. In this direction, we propose new shape functions defined as,
(b) *Local Shape Functions*:

- **DT1**: It is similar to **D2**, with an additional constraint that the time separation between two random points in reconstructed phase space is $\leq \delta$, thereby encoding only the local shape information.
- **DT2**: encodes dynamical evolution of the phase space by exponential weighting given by

$$\mathbf{DT2}_{ij} = e^{-\gamma|t_i - t_j|} * ||\mathbf{x}_i - \mathbf{x}_j||_2, \tag{5}$$

where t_i and t_j are the time indexes of the randomly selected pair of embedding vectors in the reconstructed phase space. 'δ' and 'γ' are empirically determined parameters such that $\delta, \gamma \geq 0$.

Local vs Global: The main idea behind proposing these local shape functions is that, a global shape function would consider data samples from independent repetitions (well separated in time) of a movement. Also, repetitive human movements (such as *running* and *walking*) result in trajectories which wraps around itself in reconstructed phase space, creating an artifact of having closely spaced trajectories in phase space. We believe that such an approach would not provide a robust feature representation, and we suggest the use of local shape functions instead which only considers data samples close in time.

Metric on Shape Distributions: Several metrics exist in literature to calculate the distance between histograms including chi-squared statistic (χ^2 distance), Bhattacharyya distance [49], Riemannian analysis [50] and Earth Mover's Distance (EMD) [51]. In our experiments, we provide results using Euclidean distance and chi-squared distance metrics for comparison due to their simplicity.

4.1 Test on Models

The framework was tested on the Lorenz and Rossler models to determine whether the shape feature can

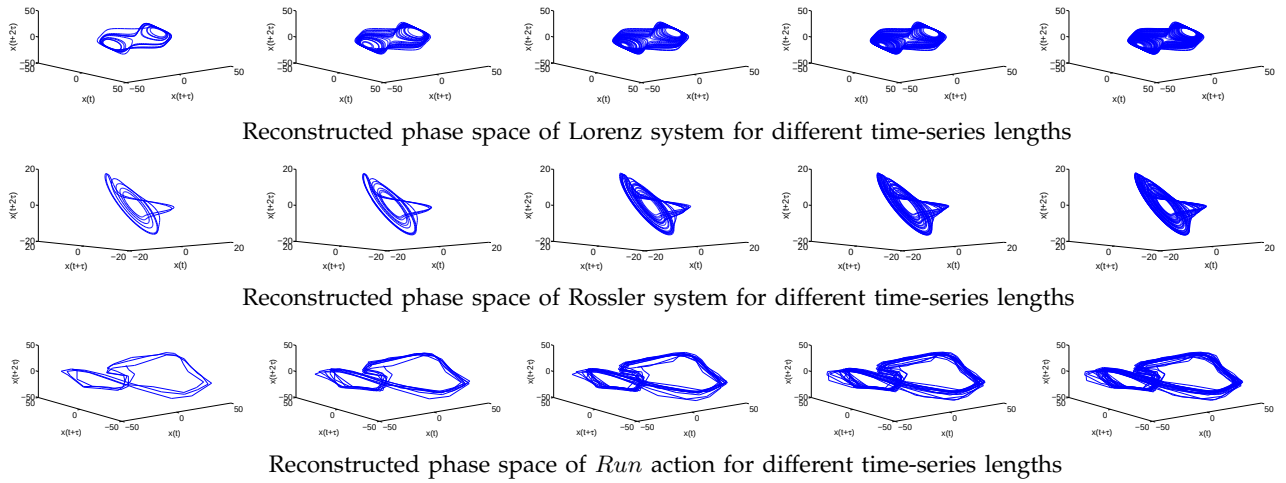

Reconstructed phase space of Lorenz system for different time-series lengths

Reconstructed phase space of Rossler system for different time-series lengths

Reconstructed phase space of *Run* action for different time-series lengths

Fig. 3: Illustration of the effect of time-series lengths on reconstructed phase space for nonlinear dynamical models like Lorenz and Rossler systems, and right-foot trajectory of a subject performing *Run* action. These examples clearly indicate that the *shape* of the reconstructed phase space does not change with time-series length, motivating feature extraction representative of the *shape* of the reconstructed phase space (as reported in Fig. 4).

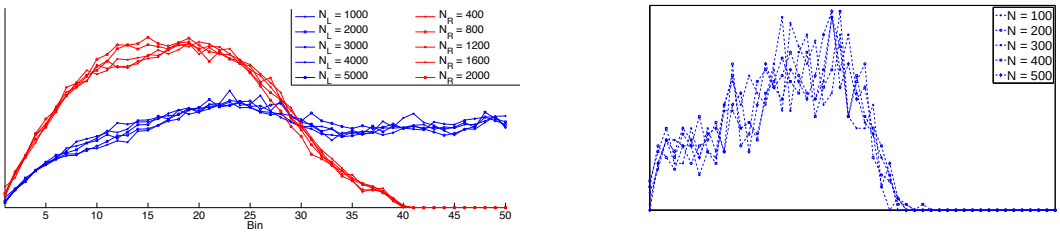

(a) Shape distribution (**D2**) of reconstructed phase space for Lorenz (blue) and Rossler (red) models for different time-series length N (N_L and N_R represent time-series lengths of Lorenz and Rossler systems respectively).

(b) Shape distribution (**D2**) of reconstructed phase space from right-foot trajectory of a subject performing *Run* action for different time-series length.

Fig. 4: Illustration of stability of the dynamical shape distribution (**D2**) extracted from reconstructed phase space for different time-series length. (a) shows the stability of **D2** distribution on Lorenz and Rossler systems while studies have reported significant error in estimation of largest Lyapunov exponent on these models (refer TABLE 1). (b) depicts the stability of **D2** distribution for trajectory data collected from right-foot of a subject performing *Run* action.

TABLE 1: Experimental results on Lorenz and Rossler models for given embedding parameters ($m_L = 3$, $\tau_L = 11$, $m_R = 3$, $\tau_R = 8$) and different time-series lengths. The true value of λ_1 for Lorenz and Rossler models are 1.50 and 0.09 respectively [14].

System	N	Calculated λ_1	% error
Lorenz	1000	1.751	16.7
	2000	1.345	-10.3
	3000	1.372	-8.5
	4000	1.392	-7.2
	5000	1.523	1.5
Rossler	400	0.0351	-61.0
	800	0.0655	-27.2
	1200	0.0918	2.0
	1600	0.0984	9.3
	2000	0.0879	-2.3

be effectively used to classify differences in shape of reconstructed phase space of nonlinear dynamical systems. We compare the performance of the proposed framework with that of largest Lyapunov exponent. The effect of time-series length on estimation of largest Lyapunov exponent was revealed by Rosenstein *et al.* [18], by evaluating the performance of the algorithm they proposed for estimation of λ_1 for various time-series lengths. The simulation results on Lorenz and Rossler models are shown in TABLE 1. Their findings indicate that the estimation error increases with reduction in time-series length (N). Fig. 3 depicts the variations in reconstructed phase space for different time-series length with defined embedding parameters. It is evident from these plots that the *shape* of the reconstructed phase space remain sufficiently similar and can be used as a discriminative feature for classification purposes. Also, from Fig. 4, the shape distribution (using **D2** shape function) was found to be stable for different time-series lengths. This striking ability of our feature representations to be robust to changes in data length will be useful in applications related to human activity analysis, where the signal

observation time is small/variable.

5 Experiments and Results

The proposed framework for representation of dynamics was evaluated on the following video-based inference tasks:
(1) Action recognition on a motion capture dataset [5].
(2) Action recognition on the MSR Action3D dataset released by Microsoft Research [52].
(3) Action quality estimation on stroke rehabilitation datasets collected in hospital and home environments [31], [53].
(4) Dynamic scene classification on the Maryland "in-the-wild" natural scene dataset [3] and the Yupenn "stabilized" scene dataset [40].

Baseline: The main contribution of our work is to propose a better way to encode dynamics compared to traditional chaotic invariants. To evaluate the effectiveness of our framework, we provide comparative results in each experiment with a feature vector [1] of traditional chaotic invariants obtained by concatenating the largest Lyapunov exponent, correlation dimension and correlation integral (for 8 values of radius) resulting in a 10-dimensional feature vector denoted as ***Chaos***. For a fair comparison, the embedding procedure is fixed as mentioned in earlier sections.

5.1 Motion Capture Dataset

In the first experiment, we evaluate the performance of the proposed framework using 3-dimensional motion capture sequences of body joints of subjects performing actions released by FutureLight, R&D division of Santa Monica Studios [5]. The dataset is a collection of five actions: *dance*, *jump*, *run*, *sit* and *walk* with $31, 14, 30, 35$ and 48 instances respectively. The classification problem on this dataset is shown to be challenging due to the presence of significant intra-class variations [5]. The data is in the form of trajectories of 3D rotation angles from 18 body joints. We use all body joints except the hip joint, to remove any effects of translational movement of the body. The 3D time-series from these 17 body joints were divided into scalar time-series resulting in a 51-dimensional vector representation for each action. Phase space reconstruction and dynamical shape feature extraction was performed. The results of the leave-one-out cross-validation approach using a nearest neighbor classifier (using Euclidean and χ^2 distance metrics) are tabulated in TABLE 2. The best classification performance we achieved was a mean accuracy of 99.37% using **DT2** dynamical shape feature, in comparison with 89.7% reported by Ali *et al.* in [5] using traditional chaotic invariants. In addition, we see that the classification performance of each dynamical shape feature is significantly better than the results achieved by using traditional chaotic invariants (***Chaos*** with $m = 3$ & $m = 5$). The proposed action modeling framework achieves near-perfect classification accuracy on the motion capture dataset even in the presence of significant intra-class variations indicating its stability. This is also evident from the examples shown in Fig. 5, where minor variations in the reconstructed phase space (in the form of intra-class variations) has not produced any significant effect on the dynamical shape feature indicating the stability of the proposed framework. From these results, we see that the dynamical shape features with temporal evolution information (**DT1** and **DT2**) performs better than the shape features **D1**, **D2** and **D3**, hence substantiating our hypothesis that shape functions with dynamical evolution information should only improve the recognition performance.

1. Code available at http://www.physik3.gwdg.de/tstool/HTML/index.html

TABLE 2: Classification rates for the various proposed dynamical shape features of phase space on the motion capture dataset with $m = 3$ (and $m = 5$ in parentheses). For comparison, we use Euclidean distance and chi-squared distance metrics as a measure of distance between probability distributions. We see that **DT2** achieves highest classification rate of 99.37%. The confusion table of the same is reported in TABLE 3.

Dynamical Shape Feature	Distance Measure	
	L_2	χ^2
Chaos	80.38 (82.28)	83.54 (85.54)
Ali *et al.*	89.70	-
D1	94.30 (94.30)	98.10 (98.10)
D2	96.84 (96.20)	96.84 (96.20)
D3	97.47 (96.84)	97.47 (97.47)
DT1	97.47 (96.20)	98.73 (98.10)
DT2	96.84 (96.20)	99.37 (99.37)

TABLE 3: Confusion table for motion capture dataset using **DT2** as the dynamical shape feature achieving mean classification rate of 99.37% when compared to 89.7% reported by Ali *et al.* in [5].

Action	Dance	Jump	Run	Sit	Walk
Dance	**30**	1	0	0	0
Jump	0	**14**	0	0	0
Run	0	0	**30**	0	0
Sit	0	0	0	**35**	0
Walk	0	0	0	0	**48**

5.2 Kinect Dataset

The framework was also evaluated on a more comprehensive dataset released by Microsoft Research called *MSR Action3D* dataset [52] having 20 action classes: *high arm wave, horizontal arm wave, hammer, hand catch, forward punch, high throw, draw x, draw tick, draw circle, hand clap, two hand wave, side boxing, bend, forward kick, side kick, jogging, tennis swing, tennis serve, golf swing, pick up & throw* with 10 subjects performing each action thrice (see Fig. 6 for example actions). The action classes in this dataset were selected to

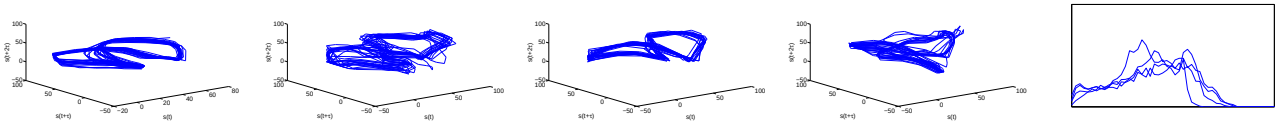

Examples of phase space reconstruction of RightLeg X-rotation time-series for 'Run' action

Shape Distribution

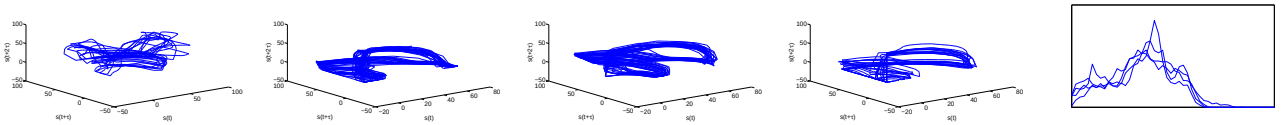

Examples of phase space reconstruction of RightLeg X-rotation time-series for 'Walk' action

Shape Distribution

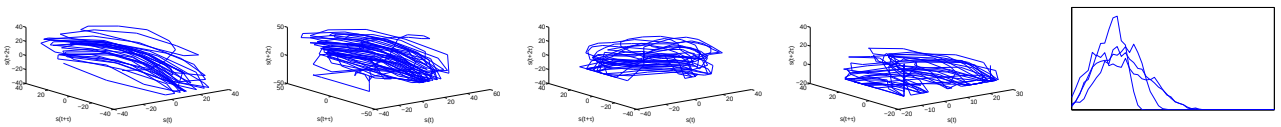

Examples of phase space reconstruction of RightLeg X-rotation time-series for 'Dance' action

Shape Distribution

Fig. 5: Illustration of the phase space reconstruction and dynamical shape feature extraction (**D2** shape feature) using four examples of *Run*, *Walk* and *Dance* action classes each from the motion capture dataset [5]. As an example, phase space reconstruction of X-rotation time-series from right leg of subjects performing these actions is shown. Embedding parameters, m was selected to be 3 and τ was calculated by method explained in section 3.4. It is evident from these examples that the 'shape' of phase space is a representative feature for an action class and can be captured using shape distributions.

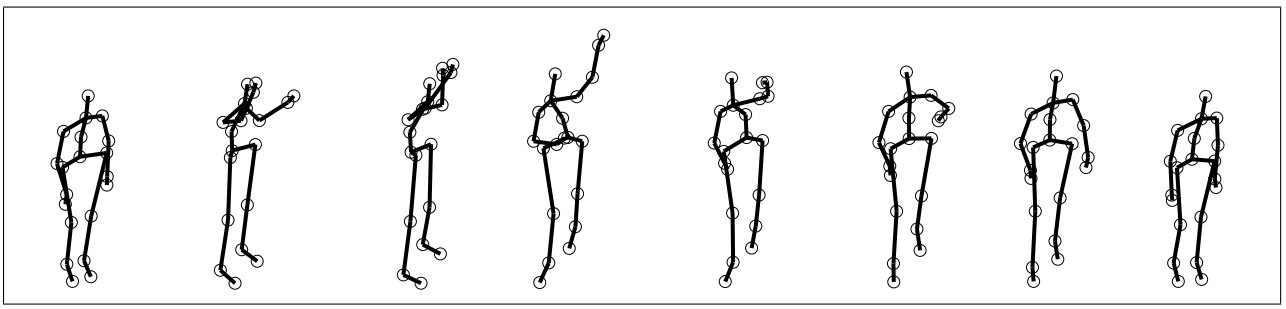

(a) **D1** (b) **D2** (c) **D3** (d) **DT1** (e) **DT2**

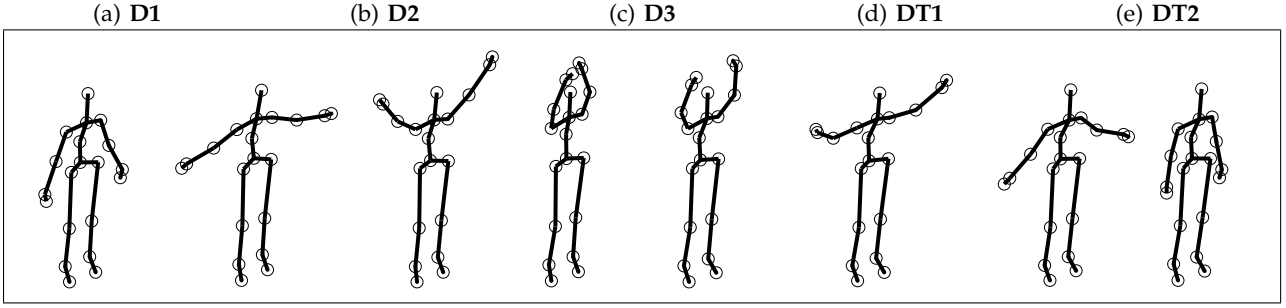

(f) **D1** (g) **D2** (h) **D3** (i) **DT1** (j) **DT2**

Fig. 6: Example actions from action class *Tennis serve* (a) and *Two hand wave* (b) from the MSR Action3D dataset. Skeleton data of 20 joints provided in the dataset will be used in our action recognition experiment. Shape distributions from reconstructed phase space using the hand trajectory from five instances each of tennis serve and two hand wave actions is shown here to illustrate the insensitivity of the framework to inter-class similarities.

TABLE 4: Classification results for cross-subject test setting where 50% subjects were used for training and the remaining 50% subjects for testing in proposed method using linear SVM with $m = 3$ (and $m = 5$ in parentheses).

	Shape Distribution					*Chaos*
Set	**D1**	**D2**	**D3**	**DT1**	**DT2**	
AS1	88.35 (86.14)	89.32 (87.13)	87.13 (86.41)	88.57 (87.38)	90.48 (89.58)	72.28 (74.56)
AS2	69.72 (63.39)	72.65 (69.75)	71.43 (72.32)	73.21 (73.50)	74.11 (70.00)	51.85 (52.40)
AS3	90.74 (84.68)	96.40 (93.69)	98.20 (96.43)	98.25 (92.92)	99.09 (96.49)	76.36 (78.86)
Avg.	**82.94 (78.07)**	**86.12 (83.52)**	**85.59 (85.05)**	**86.68 (84.60)**	**87.89 (85.34)**	66.83 (68.61)

ensure the use of arms, legs and torso by subjects to simulate interaction with gaming consoles. High similarity between classes (e.g., forward punch and hammer, high throw and pickup & throw) makes this a challenging dataset. The 20 action classes were further divided into 3 Action Sets: *AS1*, *AS2* and *AS3* in [52] to account for the large amount of computation involved in classification of these actions. The action sets 1 and 2 were intended to group actions with similar movement and action set 3 to group complex movements. The dataset provides 3D joint positions on which phase space reconstruction and extraction of shape distribution were carried out individually on every dimension (x, y & z). These shape distributions were concatenated to form our feature vector representative of any given action. The classification results on the cross-subject test setting using a linear SVM are tabulated in TABLE 4 and as seen, the proposed framework performs better than the traditional chaotic invariants. Examples shown in Fig. 6 further support our hypothesis that shape distributions can be used as discriminative feature of reconstructed phase space representative of actions. In order to illustrate the proposed framework's stability to intra-class variations and insensitivity to inter-class similarities, we compare the dynamical shape features of hand trajectory for five instances of *tennis serve* and *two hand wave* action classes. Evident from these examples is that even actions using similar hand movements are represented by dynamical shape features with enough differences to successfully recognize these actions. Furthermore, from results in TABLE 4, we see that the dynamical shape feature **DT2** has the highest overall classification accuracy, indicating that the shape distribution based on temporal evolution of phase space is better than traditional global shape representations. We have also provided classification results using a nearest neighbor classifier in TABLE 5 for a comprehensive comparison of the proposed shape distributions. Our results indicate that we achieve similar performance with both $m = 3$ and $m = 5$. In further evaluation experiments, we use $m = 3$.

5.3 Activity Quality for Stroke Rehabilitation

Our aim in this experiment is two-fold: a) to classify movements of unimpaired (neurologically normal) and impaired (stroke survivors) subjects, b) to quantitatively assess the quality of movement performed by the impaired subjects during repetitive task therapy. Fig. 7 illustrates the differences in shape of reconstructed phase space between unimpaired and impaired subjects using trajectories from the wrist marker (reflective marker placed on the subject's wrist). The experimental data was collected using a heavy marker-based system (14 markers on the right hand, arm and torso) in a hospital setting. Seven unimpaired and 15 impaired subjects perform multiple repetitions of reach and grasp movements, both on-table and elevated (the subject must move against gravity to reach the target). Each subject would perform 4 sets of reach and grasp movements to different target locations, with each set having 10 repetitions. To account for a small number of training examples, we adopt leave-one-reach-out cross validation scheme where one set of reach movement was used as testing example and rest as training examples. The stroke survivors were also evaluated by the Wolf Motor Function Test (WMFT) [33] on the day of recording, which evaluates the subject's functional ability on a scale of $1 - 5$ (with 5 being least impaired and 1 being most impaired) based on predefined functional tasks. Since our focus is on development of quantitative measures of movement quality for a home-based rehabilitation system that would use a single marker on the wrist, we only use the data corresponding to the single marker on the wrist from the heavy marker-based hospital system.

TABLE 5: Classification results for cross-subject test setting where 50% subjects were used for training and the remaining 50% subjects for testing in proposed method using nearest-neighbor classifier.

	Shape Distribution ($m = 3$)					*Chaos*	
Set	**D1**	**D2**	**D3**	**DT1**	**DT2**	$m = 3$	$m = 5$
AS1	67.00	74.62	75.73	75.05	78.43	52.30	55.67
AS2	59.63	67.66	65.77	64.47	68.21	42.53	49.23
AS3	87.83	89.96	89.66	88.11	91.13	53.45	60.59
Avg.	**71.49**	**77.41**	**77.05**	**75.87**	**79.25**	49.43	55.16

The focus of traditional methods for quantitative assessment of movement quality has been towards kinematics. Hence, in TABLE 6, we compare our results with an approach which uses kinematic analysis on the same dataset [31]. We also compare our results with the performance of traditional chaotic invariants. It is evident from these results that our framework performs better than the two promising quantitative measures for movement analysis in the field of stroke rehabilitation.

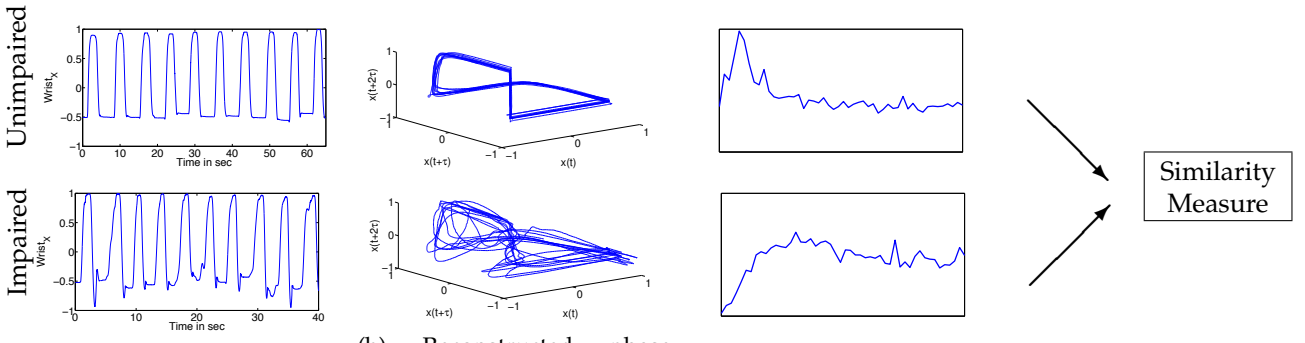

(a) Time-series data (b) Reconstructed phase space (c) Shape distribution

Fig. 7: Proposed framework for movement quality assessment and action recognition by extraction of dynamical shape feature from reconstructed phase space. (a) shows the time-series of x-location of wrist marker; its respective reconstructed phase space is shown in (b). These two exemplar trajectories are collected from the stroke rehabilitation dataset [31] and belong to unimpaired and impaired subjects respectively. The corresponding dynamical shape feature represented by shape distribution is shown in (c). Similarity measure (e.g., Euclidean distance) can be used to classify these trajectories.

TABLE 6: Comparison of classification rates for different methods using leave-one-reach-out cross-validation and nearest neighbor classifier on the stroke rehabilitation dataset.

Method	Classification Rate (%)
KIM [31]	85.2
Chaos ($m = 3$)	81.82
Chaos ($m = 5$)	83.43
D1 ($m = 3$)	**84.32**
D2 ($m = 3$)	**88.60**
D3 ($m = 3$)	**86.04**
DT1 ($m = 3$)	**87.65**
DT2 ($m = 3$)	**92.05**

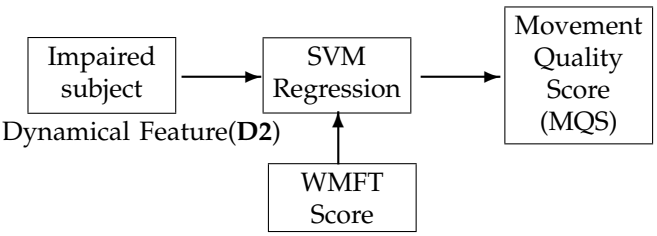

Fig. 8: Block diagram representation for learning a regressor for movement quality assessment using Functional Activity Score (FAS) from the Wolf Motor Function Test (WMFT).

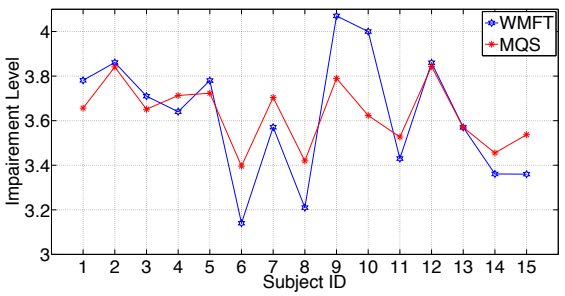

Fig. 9: Comparison between impairment level (with 5 being least impaired and 1 being most impaired) given by actual WMFT score and MQS for 15 impaired subjects. The Pearson correlation coefficient was found to be 0.8527 with a two-tail P-value of 5.35×10^{-5}, proving its statistical significance.

We also propose a framework for movement quality assessment (shown in Fig. 8) for stroke rehabilitation. Using the WMFT scores of impaired subjects, we learn a regression function using SVM to compute a movement quality score from dynamical shape feature (using **D2** shape distribution). The regressor was trained using leave-one-reach-out cross-validation technique. The outputs of the regressor were averaged per subject to get the Movement Quality Score (MQS). Fig. 9 shows a comparison between the actual WMFT score and the quality assessment score by the proposed method (MQS). The Pearson correlation coefficient between the MQS and the Function Activity Score (FAS) of the WMFT was found to be 0.8527. When we repeat the same experiment with kinematic attributes on a single wrist marker, the correlation coefficient was found to be 0.6481. In comparison, kinematic analysis of data from all 14 markers gave a correlation coefficient of 0.9041. This experiment clearly shows that the proposed framework achieves comparable results obtained by the heavy marker-based system even when using a single wrist marker, which is facilitated by the phase space reconstruction and robust feature extraction from phase space using shape distribution.

The WMFT scores are based on several functional tasks (e.g., folding a towel, picking up a pencil) and not on evaluation of the actual movements during repetitive therapy treatment (reach and grasp movements). In the above experiment, we utilize these WMFT scores as an approximate high-level quantitative measure for movement quality of impaired subjects performing reach and grasp movements, as both WMFT evaluation and 3D marker data on the wrist were obtained on the same day.

To address this conflict in collection of ground truth (movement quality labels) and trajectory data, we

TABLE 7: Comparison of performance of the proposed dynamical shape features with the performance of traditional methods used for movement quality analysis.

Method	Correlation Coefficient	MSE
KIM [31]	0.4918	0.0066
Chaos ($m = 3$)	0.4717	0.0101
Chaos ($m = 5$)	0.5089	0.0100
D1 ($m = 3$)	0.3877	0.1190
D2 ($m = 3$)	0.5029	0.0078
D3 ($m = 3$)	0.4935	0.0061
DT1 ($m = 3$)	0.4582	0.0100
DT2 ($m = 3$)	0.5510	0.0057

TABLE 8: Comparison of classification rates for various approaches on the Maryland "in-the-wild" dataset (with $m = 3$).

Class	**Chaos** [3]	**Chaos** (our) [2]	D1	D2	D3	DT1	DT2
avalanche	30	40	0	0	20	10	0
b. water	30	40	30	40	20	30	30
c. traffic	50	30	80	100	50	60	90
f. fire	30	20	10	30	30	30	30
fountain	20	0	40	30	30	30	40
i. collapse	10	0	10	0	0	10	0
landslide	10	50	0	10	20	10	20
s. traffic	20	20	20	30	30	40	30
tornado	60	10	40	70	60	50	60
v. eruption	70	0	60	70	60	40	70
waterfall	30	20	10	40	20	20	30
waves	80	40	70	80	80	90	80
whirlpool	30	20	40	50	30	70	50
Avg. (%)	36	22.31	31.54	42.31	34.62	37.69	40.77

have collected a dataset from eight stroke survivors performing reach and grasp movement tasks and have developed a rating scale for movement quality in collaboration with physical therapists. Within this scale, physical therapists would provide us an overall rating on a scale of $1 - 5$ based on the therapist's impression of the participant's performance. A score of 1 denotes that the participant could not complete the task (most impaired) and a 5 denotes that the participant performed the task with the same quality of performance as the therapist if he/she were to perform it (least impaired or unimpaired). We have collected both 3D position of the wrist and physical therapist ratings in order to make comparisons among the kinematics, our proposed measure, and the therapist ratings, across the same reach action. Utilizing the expert knowledge of the therapist ratings for these rated actions will also help us better contextualize the data to better shape our framework as a therapy tool. Using the same framework for regression as earlier, we see from TABLE 7 that the proposed framework (using **DT2**) performs better than the traditional methods for movement quality assessment in terms of correlation coefficient and mean squared error. It should be noted that the proposed framework does not require data collected from unimpaired subjects for generating MQS, while kinematic methods like KIM [31] does, making the framework more suitable to model complex tasks during therapy treatment.

2. Here "our" refers to our implementation of traditional chaotic invariants using the OpenTSTOOL package.

TABLE 9: Comparison of classification rates for various approaches on the Yupenn "stabilized" dynamic dataset (with $m = 3$).

Class	**Chaos** [3]	**Chaos** (our)[2]	D1	D2	D3	DT1	DT2
beach	27	17	77	80	77	83	77
c. street	17	70	3	87	90	100	93
elevator	40	17	7	37	10	23	17
f. fire	50	10	40	50	57	40	50
fountain	7	10	0	27	17	47	0
highway	17	17	77	47	53	33	60
l. storm	37	97	97	97	93	97	100
ocean	43	30	60	70	80	87	77
railway	3	17	60	57	23	40	60
r. river	3	87	60	90	83	87	77
sky	33	23	30	47	43	50	57
snowing	10	77	73	80	90	90	93
waterfall	10	17	50	37	30	37	37
w. farm	17	03	30	13	20	10	33
Avg. (%)	22.43	35.14	48.64	58.50	54.71	58.85	59.35

5.4 Dynamic Scene Recognition

Natural dynamic scene recognition has been gaining interest in recent years [3], [40]. In an attempt to test the generality of the proposed framework to dynamical modeling for applications in video analysis, we evaluate its performance on dynamical scene classification. In this experiment, we use the Maryland "in-the-wild" dataset [3] which is a collection of 13 classes with 10 examples per class and a larger Yupenn stabilized dynamic dataset [40] which is a collection of 14 classes with 30 examples per class. The former has videos collected from video hosting websites with no control over recording process leading to a dataset with large variations in illumination, view and scale [3]. The latter dataset was recently released to emphasize only the scene-specific temporal information rather than camera-induced ones. In addition, the scene classes in the datasets were selected to illustrate potential failure of static scene representations leading to confusion between classes (e.g., chaotic traffic and smooth traffic).

Recent research on dynamical modeling of scenes have shown that temporal (motion) information can provide better classification performance than traditional feature representations (e.g., GIST [37]) on static scenes [3], [40]. The GIST feature is based on the hypothesis that humans recognize scenes by holistic understanding of a scene [37], [54], thereby providing a global spatial representation of a scene. Shroff *et al.* employed traditional chaotic invariants to model the dynamics in the time-series of the 960-dimensional GIST descriptor extracted from each video and will be treated as our baseline. Similarly, we compare the performance of our proposed shape distribution features estimated on the 960-dimensional GIST descriptor to further support our hypothesis that proposed shape-based features can perform better than traditional chaotic invariants in video-based inference tasks.

The average classification accuracy for all the proposed dynamical shape features in comparison with traditional chaotic invariants using a nearest neighbor classifier are tabulated in TABLE 8 and 9. It is evident from these results that the proposed dynamical

shape features (**D2** and **DT2**) perform better than the traditional chaotic invariants used in literature for dynamical scene classification. Evidently it is possible to improve classification performance further by fusion of dynamical and spatial features as in [3], but here we restrict ourselves to comparison with core dynamical approaches.

6 Conclusion and Future Directions

In this paper, we have proposed a shape theoretic dynamical analysis framework for applications in action and gesture recognition, movement quality assessment for stroke rehabilitation and dynamical scene classification. We address the drawbacks of traditional measures from chaos theory for modeling the dynamics by proposing a framework combining the concepts of nonlinear time-series analysis and shape theory to extract robust and discriminative features from the reconstructed phase space. Our experiments on nonlinear dynamical models and joint trajectory data from motion capture support our hypothesis that the $shape$ of the reconstructed phase space can be used as feature representation for the above discussed applications. Furthermore, the wide range of experimental analysis on publicly available datasets for recognition of actions, gestures and scenes validate our claims. The framework was also tested on movement analysis on a finer scale, where we were interested in quantifying the *movement quality* (level of impairment) for applications in stroke rehabilitation. Our experiments using a single marker indicate that with combination of dynamical features and machine learning tools, we are able to achieve comparable performance levels to a heavy marker-based system in movement quality assessment.

In this work, we perform phase space reconstruction on every dimension independently (univariate phase space reconstruction). Our future directions will be towards employing techniques for multi-variate phase space reconstruction [55]. It has been shown in [56] that multi-variate phase space reconstruction method provides better modeling than univariate phase space reconstruction, and hence lower error in predictions for human motion. We would also like to explore the use of approximate entropy [57], a dynamical measure quantifying regularity in a time-series. The suggested number of data samples required for computation of approximate entropy is between 50 and 5000 [57], which makes it more a suitable feature representation for applications in video-based inferences.

7 Acknowledgments

This work was supported by the National Science Foundation (NSF) CAREER grant 1452163.

References

[1] V. Venkataraman, P. Turaga, N. Lehrer, M. Baran, T. Rikakis, and S. L. Wolf, "Attractor-shape for dynamical analysis of human movement: Applications in stroke rehabilitation and action recognition," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, June 2013, pp. 514–520.

[2] J. Aggarwal and M. S. Ryoo, "Human activity analysis: A review," *ACM Computing Surveys (CSUR)*, vol. 43, no. 3, p. 16, 2011.

[3] N. Shroff, P. Turaga, and R. Chellappa, "Moving vistas: Exploiting motion for describing scenes," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2010, pp. 1911–1918.

[4] N. Stergiou and L. M. Decker, "Human movement variability, nonlinear dynamics, and pathology: is there a connection?" *Human Movement Science*, vol. 30, no. 5, pp. 869–888, 2011.

[5] S. Ali, A. Basharat, and M. Shah, "Chaotic invariants for human action recognition," in *IEEE International Conference on Computer Vision*, Oct. 2007, pp. 1–8.

[6] I. N. Junejo, E. Dexter, I. Laptev, and P. Pérez, "View-independent action recognition from temporal self-similarities," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 1, pp. 172–185, 2011.

[7] M. Perc, "The dynamics of human gait," *European journal of physics*, vol. 26, no. 3, pp. 525–534, 2005.

[8] J. B. Dingwell and J. P. Cusumano, "Nonlinear time series analysis of normal and pathological human walking," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 10, no. 4, pp. 848–863, 2000.

[9] J. B. Dingwell and H. G. Kang, "Differences between local and orbital dynamic stability during human walking," *Journal of Biomechanical Engineering*, vol. 129, no. 4, pp. 586–593, 2007.

[10] R. T. Harbourne and N. Stergiou, "Movement variability and the use of nonlinear tools: principles to guide physical therapist practice," *Physical Therapy*, vol. 89, no. 3, pp. 267–282, 2009.

[11] D. J. Miller, N. Stergiou, and M. J. Kurz, "An improved surrogate method for detecting the presence of chaos in gait," *Journal of biomechanics*, vol. 39, no. 15, pp. 2873–2876, 2006.

[12] L. Ralaivola, F. d'Alché Buc *et al.*, "Dynamical modeling with kernels for nonlinear time series prediction," in *Neural Information Processing Systems*, vol. 4, 2003, pp. 129–136.

[13] A. Bissacco, A. Chiuso, Y. Ma, and S. Soatto, "Recognition of human gaits," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2001, pp. 52–57.

[14] A. Wolf, J. B. Swift, H. L. Swinney, and J. A. Vastano, "Determining lyapunov exponents from a time series," *Physica D: Nonlinear Phenomena*, vol. 16, no. 3, pp. 285–317, 1985.

[15] J.-P. Eckmann and D. Ruelle, "Ergodic theory of chaos and strange attractors," *Reviews of modern physics*, vol. 57, no. 3, pp. 617–656, 1985.

[16] M. Sano and Y. Sawada, "Measurement of the lyapunov spectrum from a chaotic time series," *Physical review letters*, vol. 55, no. 10, pp. 1082–1085, 1985.

[17] J. D. Farmer and J. J. Sidorowich, "Predicting chaotic time series," *Physical review letters*, vol. 59, no. 8, pp. 845–848, 1987.

[18] M. Rosenstein, J. Collins, and C. De Luca, "A practical method for calculating largest lyapunov exponents from small data sets," *Physica D: Nonlinear Phenomena*, vol. 65, no. 1, pp. 117–134, 1993.

[19] T. TenBroek, R. Van Emmerik, C. Hasson, and J. Hamill, "Lyapunov exponent estimation for human gait acceleration signals," *Journal of Biomechanics*, vol. 40, no. 2, p. 210, 2007.

[20] L. D. Iasemidis, D.-S. Shiau, W. Chaovalitwongse, J. C. Sackellares, P. M. Pardalos, J. C. Principe, P. R. Carney, A. Prasad, B. Veeramani, and K. Tsakalis, "Adaptive epileptic seizure prediction system," *IEEE Transactions on Biomedical Engineering*, vol. 50, no. 5, pp. 616–627, 2003.

[21] D. M. Gavrila, "The visual analysis of human movement: A survey," *Computer Vision and Image Understanding*, vol. 73, no. 1, pp. 82–98, 1999.

[22] L. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989.

[23] J. L. Casti, *Linear Dynamical Systems*. Academic Press Professional, Inc., 1986.

[24] J. Yamato, J. Ohya, and K. Ishii, "Recognizing human action in time-sequential images using hidden markov model," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 1992, pp. 379–385.

[25] A. D. Wilson and A. F. Bobick, "Learning visual behavior for gesture analysis," in *IEEE International Symposium on Computer Vision*, Nov. 1995, pp. 229–234.

[26] N. Vaswani, A. K. Roy-Chowdhury, and R. Chellappa, "Shape activity: a continuous-state hmm for moving/deforming shapes with application to abnormal activity detection," *IEEE Transactions on Image Processing*, vol. 14, no. 10, pp. 1603–1616, 2005.

[27] N. P. Cuntoor and R. Chellappa, "Epitomic representation of human activities," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2007, pp. 1–8.

[28] A. Kale, A. Sundaresan, A. Rajagopalan, N. P. Cuntoor, A. K. Roy-Chowdhury, V. Kruger, and R. Chellappa, "Identification of humans using gait," *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1163–1173, 2004.

[29] Z. Liu and S. Sarkar, "Improved gait recognition by gait dynamics normalization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 6, pp. 863–876, 2006.

[30] C. Bregler, "Learning and recognizing human dynamics in video sequences," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 1997, pp. 568–574.

[31] Y. Chen, M. Duff, N. Lehrer, H. Sundaram, J. He, S. L. Wolf, and T. Rikakis, "A computational framework for quantitative evaluation of movement during rehabilitation," in *AIP Conference Proceedings-American Institute of Physics*, vol. 1371, 2011, pp. 317–326.

[32] A. Fugl-Meyer, L. Jääskö, I. Leyman, S. Olsson, S. Steglind *et al.*, "The post-stroke hemiplegic patient. 1. a method for evaluation of physical performance." *Scandinavian journal of rehabilitation medicine*, vol. 7, no. 1, pp. 13–31, 1975.

[33] S. L. Wolf, P. A. Catlin, M. Ellis, A. L. Archer, B. Morgan, and A. Piacentino, "Assessing wolf motor function test as outcome measure for research in patients after stroke," *Stroke*, vol. 32, no. 7, pp. 1635–1639, 2001.

[34] L. Fei-Fei and P. Perona, "A bayesian hierarchical model for learning natural scene categories," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2005, pp. 524–531.

[35] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2010, pp. 3485–3492.

[36] A. Oliva and A. Torralba, "Building the gist of a scene: The role of global image features in recognition," *Progress in brain research*, vol. 155, pp. 23–36, 2006.

[37] A. Oliva and Torralba, "Modeling the shape of the scene: A holistic representation of the spatial envelope," *International Journal of Computer Vision*, vol. 42, no. 3, pp. 145–175, 2001.

[38] S. Soatto, G. Doretto, and Y. N. Wu, "Dynamic textures," in *IEEE International Conference on Computer Vision*, vol. 2, 2001, pp. 439–446.

[39] G. Doretto, A. Chiuso, Y. N. Wu, and S. Soatto, "Dynamic textures," *International Journal of Computer Vision*, vol. 51, no. 2, pp. 91–109, 2003.

[40] K. G. Derpanis, M. Lecce, K. Daniilidis, and R. P. Wildes, "Dynamic scene understanding: The role of orientation features in space and time in scene classification," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2012, pp. 1306–1313.

[41] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions," *ACM Transactions on Graphics*, vol. 21, no. 4, pp. 807–832, 2002.

[42] A. Bissacco, "Modeling and learning contact dynamics in human motion," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2005, pp. 421–428.

[43] G. P. Williams, *Chaos theory tamed*. Joseph Henry Press, 1997.

[44] H. D. Abarbanel, *Analysis of observed chaotic data*. New York: Springer-Verlag, 1996.

[45] F. Takens, "Detecting strange attractors in turbulence," *Dynamical Systems and Turbulence*, vol. 898, pp. 366–381, 1981.

[46] M. B. Kennel, R. Brown, and H. D. Abarbanel, "Determining embedding dimension for phase-space reconstruction using a geometrical construction," *Physical review A*, vol. 45, no. 6, p. 3403, 1992.

[47] M. Small, *Applied nonlinear time series analysis: applications in physics, physiology and finance*. World Scientific Publishing Company Incorporated, 2005, vol. 52.

[48] W. Tucker, "The lorenz attractor exists," *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, vol. 328, no. 12, pp. 1197–1202, 1999.

[49] A. Bhattacharyya, "On a measure of divergence between two statistical populations defined by their probability distributions," *Indian Journal of Statistics*, vol. 35, no. 99-109, p. 4, 1943.

[50] A. Srivastava, I. Jermyn, and S. Joshi, "Riemannian analysis of probability density functions with applications in vision," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2007, pp. 1–8.

[51] Y. Rubner, C. Tomasi, and L. J. Guibas, "A metric for distributions with applications to image databases," in *IEEE International Conference on Computer Vision*, Jan. 1998, pp. 59–66.

[52] W. Li, Z. Zhang, and Z. Liu, "Action recognition based on a bag of 3d points," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, Jun. 2010, pp. 9–14.

[53] M. Baran, N. Lehrer, D. Siwiak, Y. Chen, M. Duff, T. Ingalls, and T. Rikakis, "Design of a home-based adaptive mixed reality rehabilitation system for stroke survivors," in *IEEE Conference on Engineering in Medicine and Biological Society*, Aug. 2011, pp. 7602–7605.

[54] I. Biederman, "Recognition-by-components: a theory of human image understanding," *Psychological review*, vol. 94, no. 2, pp. 115–147, 1987.

[55] L. Cao, A. Mees, and K. Judd, "Dynamics from multivariate time series," *Physica D: Nonlinear Phenomena*, vol. 121, no. 1, pp. 75–88, 1998.

[56] A. Basharat and M. Shah, "Time series prediction by chaotic modeling of nonlinear dynamical systems," in *IEEE International Conference on Computer Vision*, 2009, pp. 1941–1948.

[57] S. M. Pincus, "Approximate entropy as a measure of system complexity." *Proceedings of the National Academy of Sciences*, vol. 88, no. 6, pp. 2297–2301, 1991.

Vinay Venkataraman received his M.S. degree in Electrical Engineering from Arizona State University in 2012. He is currently a doctoral student in the department of Electrical Engineering at Arizona State University. His research interests are in nonlinear dynamical analysis, computer vision and biomedical signal processing. He is a student member of IEEE.

Pavan Turaga (S'05, M'09, SM'14) is Assistant Professor in the School of Arts, Media, Engineering, and Electrical Engineering at Arizona State University. He received the B.Tech. degree in electronics and communication engineering from the Indian Institute of Technology Guwahati, India, in 2004, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park in 2008 and 2009 respectively. He then spent two years as a research associate at the Center for Automation Research, University of Maryland, College Park. His research interests are in computer vision and computational imaging with applications in activity analysis, and dynamic scene analysis, with a focus on non-Euclidean techniques for these applications. He was awarded the Distinguished Dissertation Fellowship in 2009. He was selected to participate in the Emerging Leaders in Multimedia Workshop by IBM, New York, in 2008. He received the National Science Foundation CAREER award in 2015.